arXiv:1009.1560v1 [cs.PL] 8 Sep 2010

Stack-Summarizing Control-Flow Analysis of Higher-Order Programs

Christopher Earl[1], Matthew Might[1], and David Van Horn[2]

[1] University of Utah, Salt Lake City, Utah
[2] Northeastern University, Boston, Massachusetts

Abstract. Two sinks drain precision from higher-order flow analyses: (1) merging of argument values upon procedure call and (2) merging of return values upon procedure return. To combat the loss of precision, these two sinks have been addressed independently. In the case of procedure calls, *abstract garbage collection* reduces argument merging; while in the case of procedure returns, *context-free approaches* eliminate return value merging. It is natural to expect a combined analysis could enjoy the mutually beneficial interaction between the two approaches. The central contribution of this work is a direct product of abstract garbage collection with context-free analysis. The central challenge to overcome is the conflict between the core constraint of a pushdown system and the needs of garbage collection: a pushdown system can only see the top of the stack, yet garbage collection needs to see the entire stack during a collection. To make the direct product computable, we develop "stack summaries," a method for tracking stack properties at each control state in a pushdown analysis of higher-order programs.

1 Introduction

In higher-order flow analysis [10], merging is the enemy. Merging of flow sets and control-flow paths is what destroys precision. Merging occurs in two forms: merging on call (for arguments) and merging on return (for return-flow and return values). Our goal is to alleviate argument-merging while simultaneously eliminating return-flow merging.

For an example of both kinds of merging, consider the following code:

```
(let* ((id (lambda (x) x))
       (a  (id 3))
       (b  (id 4)))
 b)
```

Flow-sensitive 0CFA makes the following inferences: (1) the two instances of the argument `x` merge together: `3` and `4`, and (2) the (implicit) continuations at applications of the identity function also merge together, causing its return values to merge in the variable `b`.

For two decades, context-sensitivity—splitting bindings, calls and returns among a finite set of abstract instances—has been the "solution" to both merging

problems. But, context-sensitivity is a finite, monotonic band-aid for an infinite, non-monotonic problem. Arguments ultimately merge because flow information accretes *monotonically*: once an analysis says that x may flow to y, it will never revoke that inference. Return-flows merge because finite flow analyses implicitly allocate a finite number of abstract stack pointers to continuations.

1.1 Two solutions

Might and Shivers developed abstract garbage collection (abstract GC) to tame the argument-merging problem [8]. Abstract GC assumes a small-step abstract interpretation [1, 2] over a finite state-space. Much like concrete GC, abstract GC finds all of the reachable addresses in an abstract heap and reclaims any unreachable addresses. With abstract GC, the abstract heap no longer grows monotonically across a small-step transition: the same abstract address has the chance to get rebound to a singleton flow set over a different value many times over, thereby making more judicious use of the abstract resources available. For programs composed of (possibly recursive) tail calls and closures which never escape, abstract garbage collection delivers perfectly precise control-flow analysis.

Pushdown control-flow analysis (PDCFA) [4], a relative of Vardoulakis and Shivers's CFA2 [13], solves the return-flow problem by using the arbitrarily large pushdown stack to model the concrete call stack; thus, continuations never merge. PDCFA can reason through arbitrary levels of recursive calls.

1.2 One problem

Our mission is to combine the benefits of both abstract garbage collection and pushdown control-flow analysis: to produce an "almost complete" control-flow analysis which eliminates *most* argument merging and *all* continuation merging. *The challenge is an apparent incompatibility between the two techniques.*

Abstract garbage collection must have the ability to search an entire state—stack included—to determine the reachable addresses. A pushdown control-flow analysis approximates the evaluation of a program, roughly speaking, as a pushdown automaton. The machine states of the PDA represent the control string, environment, and store (heap) of the evaluator; while the stack of the PDA represents the evaluator's stack, where each letter of the stack alphabet represents a continuation frame. Transitions of the PDA push and pop frames much like an abstract machine (e.g., the CESK machine) pushes and pops continuations.

When a machine like the CESK machine performs garbage collection, it crawls the stack to determine reachable heap locations. That works because the stack is explicit in each machine state: it's the K component. But in a pushdown analysis, the abstract stack is not represented in each control state. Rather, the stack's structure is scattered across the transition graph between control-states. In more detail, the data structure accumulated during pushdown analysis is a transition graph where each node contains the C, E and S components and each edge is labeled with the change to K that happens on that transition. In order to

recover the possible stack(s) at a node in the graph, the analysis must consider all the paths from the initial control state to the current state.

1.3 Our contribution: SSCFA

To complete our mission, we develop a new kind of higher-order pushdown-like control-flow analysis that includes stack *summaries* in its control states: SSCFA. To make our contribution more general, we place constraints on stack summaries (in lieu of fixing them to be reachable addresses) and we let clients supply alternate summaries, *e.g.*, all procedures live on the stack, whether the security context is privileged or unprivileged. Thus, SSCFA could drive pushdown variants of dependence analysis or even escape analysis in addition to abstract garbage collection.

The remainder of this paper is organized as follows:

- Section 2 reviews simple preliminaries for working with pushdown systems.
- Section 3 reviews pushdown control-flow analysis and Dyck state graphs.
- Section 4 introduces the problem with integrating abstract garbage collection and pushdown analysis.
- Section 5 informally introduces the notion of a *stack summary*, defines criteria for *stack summarization*, and gives example summarization strategies.
- Section 6 formally defines stack-summarizing control flow-analysis.
- Section 7 presents the computable product of stack summarizing control-flow analysis and abstract garbage collection.
- Section 8 discusses related work and Section 9 concludes.

2 Pushdown preliminaries

In this work, we make extensive use of pushdown systems. (A pushdown automaton is a specific kind of pushdown system.) There are many (equivalent) definitions of these machines in the literature, so we adapt our own definitions from [11]. Even those familiar with pushdown theory may want to skim this section to pick up our notation.

2.1 Stack actions, stack change and stack manipulation

Stacks are sequences over a alphabet Γ. Pushdown systems do much stack manipulation; to represent this more concisely, we turn stack alphabets into "action" sets where each character represents a stack change: push, pop or no change.

For each character γ in a stack alphabet Γ, the **stack-action** set $\Gamma_\pm$ contains a push (γ_+) and a pop (γ_-) character and a no-stack-change indicator (ϵ):

$$
\begin{array}{rlll}
g \in \Gamma_\pm ::= & \epsilon & & [\text{stack unchanged}] \\
\mid & \gamma_+ & \text{for each } \gamma \in \Gamma & [\text{pushed } \gamma] \\
\mid & \gamma_- & \text{for each } \gamma \in \Gamma & [\text{popped } \gamma].
\end{array}
$$

Given a string of stack actions, we can compact it into a minimal string describing net stack change. We do so through the operator $\lfloor \cdot \rfloor : \Gamma_{\pm}^{*} \to \Gamma_{\pm}^{*}$, which cancels out opposing adjacent push-pop stack actions: $\lfloor \boldsymbol{g}\ \gamma_{+}\gamma_{-}\ \boldsymbol{g}\,' \rfloor = \lfloor \boldsymbol{g}\ \boldsymbol{g}\,' \rfloor$ and $\lfloor \boldsymbol{g}\ \epsilon\ \boldsymbol{g}\,' \rfloor = \lfloor \boldsymbol{g}\ \boldsymbol{g}\,' \rfloor$, so that $\lfloor \boldsymbol{g} \rfloor = \boldsymbol{g}$, if there are no cancellations to be made in the string $\boldsymbol{\gamma}$.

2.2 Pushdown systems

A **pushdown system** is a triple $M = (Q, \Gamma, \delta)$ where Q is a finite set of control states; Γ is a stack alphabet; and $\delta \subseteq Q \times \Gamma_{\pm} \times Q$ is a transition relation. We use $\mathbb{PDS}$ to denote the class of all pushdown systems. Unlike the more widely known pushdown automaton, a pushdown system *does not recognize a language.*

For the following definitions, let $M = (Q, \Gamma, \delta)$. The **configurations** of this machine are pairs over control states and stacks: $\mathit{Configs}(M) = Q \times \Gamma^{*}$. The labeled **transition relation** $(\longmapsto_M) \subseteq \mathit{Configs}(M) \times \Gamma_{\pm} \times \mathit{Configs}(M)$ determines whether one configuration may transition to another while performing the given stack action:

$$(q, \boldsymbol{\gamma}) \longmapsto_M^{\epsilon} (q', \boldsymbol{\gamma}) \text{ iff } (q, \epsilon, q') \in \delta \qquad \text{[no change]}$$

$$(q, \gamma' : \boldsymbol{\gamma}) \longmapsto_M^{\gamma'_-} (q', \boldsymbol{\gamma}) \text{ iff } (q, \gamma'_-, q') \in \delta \qquad \text{[pop]}$$

$$(q, \boldsymbol{\gamma}) \longmapsto_M^{\gamma'_+} (q', \gamma' : \boldsymbol{\gamma}) \text{ iff } (q, \gamma'_+, q') \in \delta \qquad \text{[push].}$$

Additionally, we define:

$c \longmapsto_M c'$ iff $c \longmapsto_M^{g} c'$ for some stack action g,
$c \longmapsto_M^{\boldsymbol{g}} c'$ iff $c = c_0 \longmapsto_M^{g_1} c_1 \cdots c_{n-1} \longmapsto_M^{g_n} c_n = c'$ for some $\boldsymbol{g} = g_1 \ldots g_n$,
$c \longmapsto_M^{*} c'$ iff $c \longmapsto_M^{\boldsymbol{g}} c'$ for some $\boldsymbol{g}$.

2.3 Rooted pushdown systems

A **rooted pushdown system** is a quadruple (Q, Γ, δ, q_0) in which (Q, Γ, δ) is a pushdown system and $q_0 \in Q$ is an initial (root) state. $\mathbb{RPDS}$ is the class of all rooted pushdown systems. For a rooted pushdown system $M = (Q, \Gamma, \delta, q_0)$, we define a the **root-reachable transition relation**:

$$c \longmapsto\!\!\!\!\rightarrow_M^{g} c' \text{ iff } (q_0, \langle\rangle) \longmapsto_M^{*} c \text{ and } c \longmapsto_M^{g} c'.$$

In other words, the root-reachable transition relation also makes sure that the root control state can actually reach the transition. The root-reachable relation is overloaded to operate on control states:

$$q \longmapsto\!\!\!\!\rightarrow_M^{g} q' \text{ iff } (q, \boldsymbol{\gamma}) \longmapsto\!\!\!\!\rightarrow_M^{g} (q', \boldsymbol{\gamma}\,') \text{ for some stacks } \boldsymbol{\gamma}, \boldsymbol{\gamma}\,'.$$

3 Pushdown control-flow analysis

In this section we present the concrete and abstract semantics for the pushdown control-flow analysis (PDCFA) of a call-by-value λ-calculus, which represents the core of a higher-order programming language. To simplify presentation of the concrete and abstract semantics, we analyze programs in A-Normal Form (ANF), a syntactic discipline that enforces an order of evaluation and requires that all arguments to a function be atomic:

$$
\begin{array}{rll}
e \in \mathsf{Exp} ::= & \texttt{(let ((}v\ call\texttt{))}\ e\texttt{)} & \text{[non-tail call]} \\
\mid & call & \text{[tail call]} \\
\mid & æ & \text{[return]} \\
f, æ \in \mathsf{Atom} ::= & v \mid lam & \text{[atomic expressions]} \\
lam \in \mathsf{Lam} ::= & \texttt{(}\lambda\ \texttt{(}v\texttt{)}\ e\texttt{)} & \text{[lambda terms]} \\
call \in \mathsf{Call} ::= & \texttt{(}f\ æ\texttt{)} & \text{[applications]} \\
v \in \mathsf{Var} & \text{is a set of identifiers} & \text{[variables].}
\end{array}
$$

We use the CESK machine [5] to specify the semantics of ANF. We chose the CESK machine because it has an explicit stack. Figure 1 contains the concrete configuration-space of this machine. Each configuration contains a control-state component consisting of an expression, an environment and a store; and a continuation/stack component. Under our abstractions, the stack component of this configuration-space becomes both a finite "stack summary" in abstract control states and a stack component in the pushdown system. (See Appendix A for a review of the *finite-state* approach and comparison to the pushdown approach.)

PDCFA does not collapse the abstract stack into a finite structure like classical control-flow analysis. Instead of folding the stack into the store through frame pointers, PDCFA distributes the stack throughout an enriched abstract transition system. The abstract configuration-space of pushdown control-flow analysis (Figure 1) is similar to concrete formulation.

3.1 Concrete semantics and PDCFA

Next, we define the concrete semantics of ANF and pushdown control-flow analysis simultaneously. Specifically, we define program-to-machine injection, atomic expression evaluation, reachable configurations/control states, the transition relation and a resource-allocation parameter. The abstraction functions that connect the concrete configuration-space to the abstract configuration-space are straightforward structural abstraction functions. (Formal definitions of these abstractions can be found in Appendix B.)

Program injection The concrete program-injection function pairs an expression with an empty environment, store and stack to create the initial configuration:

$$c_0 = \mathcal{I}(e) = (e, [], [], \langle\rangle).$$

$c \in \mathit{Conf} = \mathit{State} \times \mathit{Kont}$	$\hat{c} \in \widehat{\mathit{Conf}} = \widehat{\mathit{State}} \times \widehat{\mathit{Kont}}$	[configurations]
$\varsigma \in \mathit{State} = \mathsf{Exp} \times \mathit{Env} \times \mathit{Store}$	$\hat{\varsigma} \in \widehat{\mathit{State}} = \mathsf{Exp} \times \widehat{\mathit{Env}} \times \widehat{\mathit{Store}}$	[states]
$\rho \in \mathit{Env} = \mathsf{Var} \rightharpoonup \mathit{Addr}$	$\hat{\rho} \in \widehat{\mathit{Env}} = \mathsf{Var} \rightharpoonup \widehat{\mathit{Addr}}$	[environments]
$\sigma \in \mathit{Store} = \mathit{Addr} \to \mathit{Clo}$	$\hat{\sigma} \in \widehat{\mathit{Store}} = \widehat{\mathit{Addr}} \to \mathcal{P}\left(\widehat{\mathit{Clo}}\right)$	[stores]
$\mathit{clo} \in \mathit{Clo} = \mathsf{Lam} \times \mathit{Env}$	$\widehat{\mathit{clo}} \in \widehat{\mathit{Clo}} = \mathsf{Lam} \times \widehat{\mathit{Env}}$	[closures]
$\kappa \in \mathit{Kont} = \mathit{Frame}^*$	$\hat{\kappa} \in \widehat{\mathit{Kont}} = \widehat{\mathit{Frame}}^*$	[stacks]
$\phi \in \mathit{Frame} = \mathsf{Var} \times \mathsf{Exp} \times \mathit{Env}$	$\hat{\phi} \in \widehat{\mathit{Frame}} = \mathsf{Var} \times \mathsf{Exp} \times \widehat{\mathit{Env}}$	[stack frames]
$a \in \mathit{Addr}$ is an infinite set	$\hat{a} \in \widehat{\mathit{Addr}}$ is a *finite* set	[addresses]

Fig. 1. Configuration-space for CESK machine and pushdown control-flow analysis.

We define two abstract injection functions—one that produces an initial abstract control state, and one that produces an initial abstract configuration. The control-state injector $\hat{\mathcal{I}}_{\hat{\varsigma}} : \mathsf{Exp} \to \widehat{\mathit{State}}$ pairs an expression with an empty environment and store to create the initial abstract state:

$$\hat{\varsigma}_0 = \hat{\mathcal{I}}_{\hat{\varsigma}}(e) = (e, [], []).$$

The configuration injector $\hat{\mathcal{I}}_{\hat{c}} : \mathsf{Exp} \to \widehat{\mathit{Conf}}$ tacks on an empty stack:

$$\hat{c}_0 = \hat{\mathcal{I}}_{\hat{c}}(e) = (\hat{\mathcal{I}}_{\hat{\varsigma}}(e), \langle\rangle).$$

Atomic expression evaluation The atomic expression evaluator, $\mathcal{A} : \mathsf{Atom} \times \mathit{Env} \times \mathit{Store} \rightharpoonup \mathit{Clo}$ (or, $\hat{\mathcal{A}} : \mathsf{Atom} \times \widehat{\mathit{Env}} \times \widehat{\mathit{Store}} \to \mathcal{P}(\widehat{\mathit{Clo}})$ in the abstract), returns the value of an atomic expression in the context of an environment and a store:

$$\mathcal{A}(\mathit{lam}, \rho, \sigma) = (\mathit{lam}, \rho) \qquad \hat{\mathcal{A}}(\mathit{lam}, \hat{\rho}, \hat{\sigma}) = \{(\mathit{lam}, \rho)\} \qquad \text{[closure creation]}$$
$$\mathcal{A}(v, \rho, \sigma) = \sigma(\rho(v)) \qquad \hat{\mathcal{A}}(v, \hat{\rho}, \hat{\sigma}) = \hat{\sigma}(\hat{\rho}(v)) \qquad \text{[variable look-up].}$$

Reachable configurations The program evaluator $\mathcal{E} : \mathsf{Exp} \to \mathcal{P}(\mathit{Conf})$ (or, $\hat{\mathcal{E}} : \mathsf{Exp} \to \mathcal{P}(\widehat{\mathit{Conf}})$ in the abstract) computes all of the configurations reachable from the initial configuration:

$$\mathcal{E}(e) = \{c : \mathcal{I}(e) \Rightarrow^* c\} \qquad \hat{\mathcal{E}}(e) = \left\{\hat{c} : \hat{\mathcal{I}}_{\hat{c}}(e) \rightharpoonup^* \hat{c}\right\}.$$

Since the stack's depth is unbounded, the number of reachable configurations in both the concrete *and* abstract semantics could be infinite.

Transition relation The concrete transition, $c \Rightarrow c'$, and its abstract counterpart, $\hat{c} \rightharpoonup \hat{c}'$, each have three rules. The first rule handles tail calls by evaluating the

function into a closure, evaluating the argument into a value and then moving to the body of the λ-term within the closure:

$$c = (\overbrace{([\![(f\ ae)]\!], \rho, \sigma)}^{\varsigma}, \kappa) \Rightarrow ((e, \rho'', \sigma'), \kappa), \text{ where}$$
$$([\![(\lambda\ (v)\ e)]\!], \rho') = \mathcal{A}(f, \rho, \sigma) \qquad \rho'' = \rho'[v \mapsto a]$$
$$a = alloc(v, \varsigma) \qquad \sigma' = \sigma[a \mapsto \mathcal{A}(ae, \rho, \sigma)]$$

$$\hat{c} = (\overbrace{([\![(f\ ae)]\!], \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, \hat{\kappa}) \rightharpoonup ((e, \hat{\rho}'', \hat{\sigma}'), \hat{\kappa}), \text{ where}$$
$$([\![(\lambda\ (v)\ e)]\!], \hat{\rho}') \in \hat{\mathcal{A}}(f, \hat{\rho}, \hat{\sigma}) \qquad \hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}]$$
$$\hat{a} = \widehat{alloc}(v, \hat{\varsigma}) \qquad \hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathcal{A}}(ae, \hat{\rho}, \hat{\sigma})].$$

In the abstract semantics, the tail-call transition is nondeterministic, since multiple abstract closures may be invoked.

A non-tail call builds a frame, adds it to the stack, and evaluates the call:

$$(([\![\texttt{(let ((}v\ call\texttt{)) }e\texttt{)}]\!], \rho, \sigma), \kappa) \Rightarrow ((call, \rho, \sigma), (v, e, \rho) : \kappa)$$
$$(([\![\texttt{(let ((}v\ call\texttt{)) }e\texttt{)}]\!], \hat{\rho}, \hat{\sigma}), \hat{\kappa}) \rightharpoonup ((call, \hat{\rho}, \hat{\sigma}), (v, e, \hat{\rho}) : \hat{\kappa}).$$

A function return pops the top frame of the stack and uses that frame to continue the computation after binding the return value to the frame's variable:

$$c = (\overbrace{(ae, \rho, \sigma)}^{\varsigma}, (v, e, \rho') : \kappa) \Rightarrow ((e, \rho'', \sigma'), \kappa), \text{ where}$$
$$a = alloc(v, \varsigma) \qquad \rho'' = \rho'[v \mapsto a] \qquad \sigma' = \sigma[a \mapsto \mathcal{A}(ae, \rho, \sigma)]$$

$$\hat{c} = (\overbrace{(ae, \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, (v, e, \hat{\rho}') : \hat{\kappa}') \rightharpoonup ((e, \hat{\rho}'', \hat{\sigma}'), \hat{\kappa}'), \text{ where}$$
$$\hat{a} = \widehat{alloc}(v, \hat{\varsigma}) \qquad \hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}] \qquad \hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathcal{A}}(ae, \hat{\rho}, \hat{\sigma})].$$

Allocation, polyvariance and context-sensitivity The address-allocation function is an opaque parameter in both semantics. For the concrete semantics, letting addresses be natural numbers suffices, and then the allocator can use the lowest unused address: $Addr = \mathbb{N}$ and $alloc(v, (e, \rho, \sigma, \kappa)) = 1 + \max(dom(\sigma))$. The opacity is useful because abstract semantics also parameterize allocation—to provide a knob to tune the polyvariance and context-sensitivity of the resulting analysis—and allowing the abstract semantics to choose a particular concrete allocation function can simplify proofs of soundness.

3.2 Removing the explicit stack

The reachable subset of the abstract configuration-space for any program could be infinite. (There is no bound on the depth of the stack, so there are an infinite

number of stacks and therefore an infinte number of configurations.) Consequently, the naïve exploration of the reachable abstract configurations used in classical flow analyses may not terminate. Fortunately, because the abstract semantics describe a pushdown system, we can construct a finite (computable) description of the reachable configurations. Specifically, we can construct a labeled transition system in which nodes are control states, and labels on edges denote stack change.

We define the legal transitions in any such graph through the transition relation $(\hat{\curvearrowright}) \subseteq \widehat{State} \times \widehat{Frame}_{\pm} \times \widehat{State}$. Three rules define this relation: one determining when to push, one when to pop and the last when to leave the stack unchanged. The labels on each transition are the stack action for the transition:

$$\hat{\varsigma} \overset{\epsilon}{\hat{\curvearrowright}} \hat{\varsigma}' \text{ iff } \hat{c} = (\hat{\varsigma}, \hat{\kappa}) \rightharpoonup (\hat{\varsigma}', \hat{\kappa}) = \hat{c}', \text{ for any stack } \hat{\kappa} \qquad \text{[tail call]}$$

$$\hat{\varsigma} \overset{\hat{\phi}_+}{\hat{\curvearrowright}} \hat{\varsigma}' \text{ iff } \hat{c} = (\hat{\varsigma}, \hat{\kappa}) \rightharpoonup (\hat{\varsigma}', \hat{\phi} : \hat{\kappa}) = \hat{c}', \text{ for any stack } \hat{\kappa} \qquad \text{[non-tail call]}$$

$$\hat{\varsigma} \overset{\hat{\phi}_-}{\hat{\curvearrowright}} \hat{\varsigma}' \text{ iff } \hat{c} = (\hat{\varsigma}, \hat{\phi} : \hat{\kappa}) \rightharpoonup (\hat{\varsigma}', \hat{\kappa}) = \hat{c}', \text{ for any stack } \hat{\kappa} \qquad \text{[return].}$$

From this transition relation, we build a rooted pushdown system (Q, Γ, δ, q_0) for a program e such that $Q = \widehat{States}$, $\Gamma = \widehat{Frame}$, $\delta = (\hat{\curvearrowright})$, and $q_0 = \hat{\mathcal{I}}_{\hat{\varsigma}}(e)$. The subset of this rooted pushdown system reachable over legal paths provides a finite description of the original configuration-space. This finite subset is a **Dyck state graph** (DSG) [4]. (A path is legal only if all of the pops match up with pushes; there can be unmatched pushes left over.) Several techniques can compute the Dyck state graph; for an efficient technique specific to PDCFA, we defer to our recent work [4] or the algorithm as modified in Appendix C.

4 Adding abstract garbage collection to PDCFA

In the classical version of abstract garbage collection, the abstract interpretation "collects" each configuration before each transition [8]. To collect a state, it explores the state to find the reachable abstract addresses, and then it discards unreachable addresses from the store, *i.e.*, it maps them to the empty set.

Suppose we were to add abstract garbage collection to PDCFA. At first, we might try collecting a control state prior to adding an edge. But, this approach doesn't work: to know the reachable addresses of a configuration, the analysis must have access to the stack paired with the control state. Unfortunately, the stack has been distributed across the Dyck state graph being accreted during the analysis. To determine the possible stacks paired with a control state, the analysis must consider all legal paths to that control state.

Considering all possible paths to a control state is expensive, and troublesome in any event, since there could be an infinite number of such paths. A better solution would allow the analysis to iteratively compute properties of stacks—like reachable addresses—and store these summaries at individual control states. We call this solution stack summaries.

5 Stack summaries

As PDCFA constructs a DSG, it accretes reachable control states one edge at a time. Each time it adds a labeled edge, it is abstractly executing the transition relation ($\rightharpoonup$). To perform abstract garbage collection before each transition, the analysis must know the reachable addresses for all configurations described by paths to that state. To accomplish this, we add a stack summary to each control state. A stack summary is a client-defined finite abstraction of a stack. To perform abstract garbagae collection, we will instantiate this summary to be the reachable addresses in the stack.

A stack summary describes some property of the stack, *e.g.*, the topmost frame, the reachable addresses, the privilege level of the current context. With respect to our analysis, the set $\widehat{Summary}$ is a parameter containing all stack summaries, and we denote an individual stack summary as $\widehat{ss}$. A summarizing function, $\alpha_S : Stack \to \widehat{Summary}$, walks a stack to compute a summary. Every stack summary regime also requires a push function parameter, $\mathbf{push} : \widehat{Frame} \times \widehat{Summary} \to \widehat{Summary}$, which computes the abstract effect of pushing a frame on a summary. There are three requirements on stack summaries:

1. Summaries must be able to represent all possible stacks.
2. The set of summaries must be finite.
3. Summaries must form a lattice ($\sqsubseteq_S$).

In addition, the push function must faithfully simulate concrete push; formally:

$$\text{if } \alpha(\phi) \sqsubseteq \hat{\phi} \text{ and } \alpha_S(\kappa) \sqsubseteq_S \widehat{ss}, \text{ then } \alpha_S(\phi : \kappa) \sqsubseteq \mathbf{push}(\hat{\phi}, \widehat{ss}).$$

We can efficiently percolate stack summaries through the construction of a Dyck state graph, so that the algorithm never has to reconsider all paths to a control state. In fact, the algorithm never considers an entire path all at once; it propagates summaries edge-by-edge. To extend the Dyck-state-graph-construction algorithm, we need to consider three cases: what is the effect of adding a push; what is the effect of adding a pop; and what is the effect of a stack no-op? In this section, we describe the core of the algorithm informally but with sufficient detail to motivate the high-level idea. In the next section, we'll describe the system-space of the algorithm formally, and Appendix C contains the algorithm for computing a Dyck state graph with stack summaries.

5.1 Propagating stack summaries during DSG construction

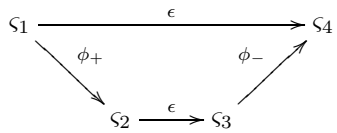
The propagation of stack summaries across no-op and pop edges during DSG construction is agnostic of the particular stack summary in use. Propagation across push edges is fully factored into the push function parameter, **push**.

The summarizing no-op operation When the Dyck state graph construction algorithm needs to propagate summaries across an edge which does not change the stack, the new stack summary is identical to the old sumary: when there is no stack change, there is no change to stack summaries.

The summarizing pop operation The pop operation, like the no-op operation, can be handled without knowledge of the particular stack summary in use. In PDCFA, every pop transition has at least one matching push transition. The stack summaries after a pop are those stack summaries that can reach the new state with no net stack change.

These states are easy to find, because the DSG construction algorithm maintains an ϵ-closure graph in addition to the control-state transition graph. Edges in the ϵ-closure graph connect states reachable through no net stack change.

Diagramatically, we know that the stack summary at state ς_4 in the following is the same as the stack summary for ς_1:



The summarizing push operation Pushing a frame onto a stack makes a local change to the stack. However, pushing a frame onto a stack may nontrivially change the summary. The operation **push** must be able to determine the new stack summary, so that when a push edge is introduced, the **push** operation determines the subsequent summary.

5.2 Example: A frame-set summary

The frame-set summary is both general and useful. The frame-set summary is the set of (abstract) frames currently in the stack:

$$\widehat{Summary}_{fs} = \mathcal{P}\left(\widehat{Frame}\right).$$

This summary ignores order and repetition in favor of finite size and a simple (subset-based) lattice:

$$\widehat{ss} \sqsubseteq_S^{fs} \widehat{ss}' \text{ iff } \widehat{ss} \subseteq \widehat{ss}'.$$

The summarization function for the frame set summary, $\alpha_S^{fs} : \widehat{Stack} \to \widehat{Summary}_{fs}$, abstracts each frame and keeps it in a set:

$$\alpha_S^{fs}\langle\phi_1, \ldots, \phi_n\rangle = \{\alpha_{Frame}(\phi_1), \ldots, \alpha_{Frame}(\phi_n)\}.$$

The push operation $\mathbf{push}^{fs} : \widehat{Frame} \times \widehat{Summary}_{fs} \to \widehat{Summary}_{fs}$ simply adds the new frame to the set: $\mathbf{push}^{fs}(\hat{\phi}, \widehat{ss}) = \{\hat{\phi}\} \cup \widehat{ss}$.

5.3 Example: A reachable-addresses summary

The reachable-addresses summary is the set of all the addresses directly touchable by a frame on the stack. We formally define *touch* through the touch function, $\mathcal{T}_f : \widehat{Frame} \to \widehat{Addr}$, which returns the addresses within the given frame:

$$\mathcal{T}_f(v, e, \hat{\rho}) = \{\hat{\rho}(v') : v' \in \mathit{free}(e) - \{v\}\}.$$

The summary-space is the set of addresses:

$$\widehat{Summary}_{ra} = \mathcal{P}\left(\widehat{Addr}\right).$$

The order on summaries is subset inclusion:

$$\widehat{ss} \sqsubseteq_S^{ra} \widehat{ss}' \text{ iff } \widehat{ss} \subseteq \widehat{ss}'.$$

The reachable address summarization function, $\alpha_S^{ra} : \widehat{Stack} \to \widehat{Summary}_{ra}$, finds the reachable addresses of each abstracted frame and keeps them in a set:

$$\alpha_S^{ra}\langle\phi_1, \ldots, \phi_n\rangle = \mathcal{T}_f(\alpha_{Frame}(\phi_1)) \cup \cdots \cup \mathcal{T}_f(\alpha_{Frame}(\phi_n)).$$

The push operation $\mathbf{push}^{ra} : \widehat{Frame} \times \widehat{Summary}_{ra} \to \widehat{Summary}_{ra}$ adds the reachable addresses from the new frame to the set:

$$\mathbf{push}^{ra}(\hat{\phi}, \widehat{ss}) = \mathcal{T}_f(\hat{\phi}) \cup \widehat{ss}.$$

The reachable address summary provides the information about the stack needed for abstract garbage collection with pushdown control-flow analysis.

6 SSCFA: Stack-summarizing control-flow analysis

In the last section, we defined stack summaries and motivated their implementation informally. In this section, we formally define the configuration-space and an abstract pushdown semantics for stack-summarizing control-flow analysis (SSCFA). Appendix C describes a formal algorithm for creating a finite model of the reachable state-space for SSCFA.

6.1 Abstract configuration-space

The only change between the configuration-spaces for the pushdown control-flow analysis and the stack-summarizing control-flow analysis is that configurations contain stack summaries instead of stacks:

$$\hat{c} \in \widehat{Conf} = \widehat{State} \times \widehat{Summary} \qquad \text{[configurations]}.$$

6.2 Abstract pushdown semantics

The abstract transition relation for SSCFA is similar to the transition relation for PDCFA. The transition relation, $(\approx\!\!>) \subseteq \widehat{Conf} \times \widehat{Frame}_{\pm} \times \widehat{Conf}$ has three rules. With respect to a program e, we can define a rooted pushdown system, $M_{\text{SS}} = (\widehat{Conf}, \widehat{Frame}, (\approx\!\!>), \hat{c}_0)$, where $\hat{c}_0 = (e, [], [], \bot_{SS})$.

A tail call leaves the stack unchanged:

$$\overbrace{((\llbracket (f\ \text{æ})\rrbracket, \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, \widehat{ss}) \overset{\epsilon}{\approx\!>} ((e, \hat{\rho}'', \hat{\sigma}'), \widehat{ss}), \text{ where}$$
$$(\llbracket (\lambda\ (v)\ e)\rrbracket, \hat{\rho}') \in \hat{\mathcal{A}}(f, \hat{\rho}, \hat{\sigma}) \qquad \hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}]$$
$$\hat{a} = \widehat{alloc}(v, \hat{\varsigma}) \qquad \hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathcal{A}}(\text{æ}, \hat{\rho}, \hat{\sigma})].$$

A non-tail call builds a frame, adds it to the summary, and evaluates the call:

$$((\llbracket \texttt{(let ((}v\ call\texttt{))}\ e\texttt{)}\rrbracket, \hat{\rho}, \hat{\sigma}), \widehat{ss}) \overset{\hat{\phi}_+}{\approx\!>} ((call, \hat{\rho}, \hat{\sigma}), \widehat{ss}'), \text{ where}$$
$$\hat{\phi} = (v, e, \hat{\rho}) \qquad \widehat{ss}' = \mathbf{push}(\hat{\phi}, \widehat{ss}).$$

A function return pops the top frame off the stack. It also restores older stack summaries. Thus, the algorithm must know all of the abstract configurations on paths from the initial configuration that can reach the current configuration on a path whose net stack change is the the frame to be popped; we find these abstract configurations using the $\mathbf{pred} : \widehat{Conf} \times \widehat{Frame} \to \mathcal{P}\left(\widehat{Conf}\right)$ function:

$$\mathbf{pred}(\hat{c}, \hat{\phi}) = \left\{ \hat{c}' : \hat{c} \longmapsto\!\!\!\!\!\rightarrow^{\boldsymbol{\phi'}}_{M_{\mathrm{SS}}} \hat{c}' \text{ and } \lfloor \boldsymbol{\phi'} \rfloor = \phi_+ \right\}.$$

The transition rule for pop is then straightforward:

$$\overbrace{((\text{æ}, \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, \widehat{ss}) \overset{\hat{\phi}_-}{\approx\!>} ((e, \hat{\rho}'', \hat{\sigma}'), \widehat{ss}'), \text{ where}$$
$$(_, \widehat{ss}') \in \mathbf{pred}(\hat{c}, \hat{\phi}) \qquad \hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}]$$
$$(v, e, \hat{\rho}') = \hat{\phi} \qquad \hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathcal{A}}(\text{æ}, \hat{\rho}, \hat{\sigma})]$$
$$\hat{a} = \widehat{alloc}(v, \hat{\varsigma}).$$

6.3 Soundness of stack-summarizing control-flow analysis

The soundness of Dyck state graphs has been proved in [4]. However, the soundness of the stack summaries is provided below for the first time:

Theorem 1. *If $\alpha(c) \sqsubseteq \hat{c}$, $c \Rightarrow c'$ and $c_0 \Rightarrow^* c$, then there exists $\hat{c}' \in \widehat{Conf}$ such that $\alpha(c') \sqsubseteq \hat{c}'$ and $\hat{c} \approx\!> \hat{c}'$.*

Proof (sketch). Let $c = (\varsigma, \kappa)$, $c' = (\varsigma', \kappa')$ and $\hat{c} = (\hat{\varsigma}, \widehat{ss})$, such that $\alpha(c) \sqsubseteq \hat{c}$. We know by theorems in [4] that there exists a state $\hat{\varsigma}'' \in \widehat{State}$ such that $\alpha(\varsigma') \sqsubseteq \hat{\varsigma}''$. We also know that the first stack is subsumed by the first stack summary: $\alpha_S(\kappa) \sqsubseteq_S \widehat{ss}$. So we must prove that there exists a stack summary $\widehat{ss}' \in \widehat{Summary}$ such that $\alpha_S(\kappa') \sqsubseteq_S \widehat{ss}'$ and $\hat{c} \approx\!> (\hat{\varsigma}'', \widehat{ss}')$. The proof continues with a case-wise analysis on the type of the transition as well as strong induction based upon the length of the path to configuration c. See Appendix D for details.

7 SSCFA with Abstract Garbage Collection

Having constructed a framework for iteratively synthesizing stack summaries during computation of a finite model for a pushdown system, we can integrate abstract garbage collection. In this section, we assume the "reachable addresses" stack summary is in use. We term this analysis SSΓCFA, the product of stack-summarizing control-flow analysis and abstract garbage collection (also called ΓCFA). SSΓCFA is a "best of both worlds" combination: it has all the argument precision advantages of abstract garbage collection and all the return-flow precision advantages of PDCFA.

As with classical abstract garbage collection, we must define what makes an address or value reachable. Essentially, an object is *reachable* if it may be used either in the current configuration or in a subsequent configuration. If an address is reachable, all the values bound to it are also reachable. Values (closures and frames) reference addresses through their environments, which are reachable as well. Because values touch addresses and addresses touch values, finding reachable addresses and values amounts to a bipartite graph search. The concrete values of unreachable addresses will never be used again during the course of the computation; thus, it is safe to set the values of these addresses to bottom within the store.

The reachability exploration of the store begins with the addresses that the current configuration $\hat{c}$ can immediately reach, called the root set, $\mathit{Root}(\hat{c})$. The root function, $\mathit{Root} : \widehat{\mathit{Conf}} \to \mathcal{P}\left(\widehat{\mathit{Addr}}\right)$, returns the root set for a configuration:

$$\mathit{Root}((e, \hat{\rho}, \hat{\sigma}), \widehat{ss}) = \widehat{ss} \cup \{\hat{\rho}(v) : v \in \mathit{free}(e)\},$$

where the function $\mathit{free} : \mathsf{Exp} \to \mathcal{P}(\mathsf{Var})$ returns the free variables in the given expression. The root set contains all the addresses bound to free variables in the expression, e, as well as the addresses in the reachable address summary.

The touch function, $\mathcal{T}_c : \widehat{\mathit{Clo}} \to \widehat{\mathit{Addr}}$, finds addresses referenced in closures: $\mathcal{T}_c(\mathit{lam}, \hat{\rho}) = \{\hat{\rho}(v) : v \in \mathit{free}(\mathit{lam})\}$. The touching relation, $(\rightsquigarrow_{\mathcal{T},\hat{\sigma}}) : \widehat{\mathit{Addr}} \to \widehat{\mathit{Addr}}$ links addresses directly to addresses:

$$\hat{a} \rightsquigarrow_{\mathcal{T},\hat{\sigma}} \hat{a}' \text{ iff } \hat{a}' \in \mathcal{T}_c(\widehat{\mathit{val}}) \text{ and } \widehat{\mathit{val}} \in \hat{\sigma}(\hat{a}).$$

With this relation, finding all reachable addresses of a configuration $\hat{c}$ becomes the transitive closure of the touching relation:

$$\mathcal{R}(\hat{c}) = \left\{\hat{a}' : \hat{a} \rightsquigarrow^*_{\mathcal{T},\hat{\sigma}} \hat{a}' \text{ and } \hat{a} \in \mathit{Root}(\hat{c})\right\}.$$

Finally, we define the abstract garbage collector itself, $\mathit{AGC} : \widehat{\mathit{Conf}} \to \widehat{\mathit{Conf}}$, which simply restricts the store to the reachable addresses:[3]

$$\mathit{AGC}(\hat{c}) = (e, \hat{\rho}, \hat{\sigma}|\mathcal{R}(\hat{c}), \widehat{ss}), \text{ where } \hat{c} = (e, \hat{\rho}, \hat{\sigma}, \widehat{ss}).$$

[3] We define function restriction, $f|X$, so that $f|X = \lambda x.\textbf{if } x \in X \textbf{ then } f(x) \textbf{ else } \bot$.

The abstract transition relation for SSΓCFA, $(\hat{\approx}\!\!\gg_{AGC}) \subseteq \widehat{Conf} \times \widehat{Conf}$, needed for stack-summarizing control-flow analysis with abstract garbage collection extends the abstract transition relation to collect before each transition:

$$\hat{c} \hat{\approx}\!\!\gg_{AGC} \hat{c}' \text{ iff } AGC(\hat{c}) \hat{\approx}\!\!\gg \hat{c}'.$$

The soundness theorems and their proofs for classical abstract garbage collection are in Chapter 6 of [7]; they adapt readily to our pushdown framework.

8 Related Work

Stack summarization, the central contribution of this paper, overcomes the apparent incompatibilities of two orthogonal anti-merging techniques designed to improve precision: abstract garbage collection [8] and pushdown control-flow analysis [4, 13]. As such, this work directly builds upon both techniques, as well as classical control-flow analysis [10], abstract machines [5], and abstract interpretation [1, 2] in general.

Abstract garbage collection [8, 12] curbs argument-merging, but it has not yet been applied to anything beyond classical control-flow analysis.

Vardoulakis and Shivers's CFA2 [13] is the precursor to the pushdown control-flow analysis [4] presented in Section 3. CFA2 is a table-driven summarization algorithm that exploits the balanced nature of calls and returns to improve return-flow precision in a control-flow analysis. While CFA2 uses a concept called "summarization," it is a summarization of execution paths of the analysis, roughly equivalent to Dyck state graphs rather than our stack summaries.

In terms of recovering precision, pushdown control-flow analysis [4] is the dual to abstract garbage collection: it focuses on the global interactions of configurations via transitions to precisely match push-pop/call-return, thereby eliminating all return-flow merging. However, pushdown control-flow analysis does nothing to improve argument merging.

In the context of first-order languages, pushdown approaches to analysis are well-established. Reps *et al.* [9] uses a summarization algorithm to compute a Dyck-state-graph-like solution. Debray and Proebsting [3] develop an analysis with perfect return-flow in the presence of tail calls. For higher-order languages, finite-state approaches *approximating* the pushdown precision of return-flow have been explored by Midtgaard and Jensen [6] and Van Horn and Might [12]. Our work extends the pushdown approach to higher-order languages with tail calls, and produces stack summaries to enable abstract garbage collection.

9 Conclusion

We presented SSΓCFA, a synergistic fusion of pushdown analysis and abstract garbage collection to combat the twin sinks for precision in higher-order flow analysis: merging in arguments, and merging in return-flow. In order to create SSΓCFA, we had to first create SSCFA, a pushdown control-flow analysis for

higher-order programs capable of iteratively synthesizing summaries of stack properties; in this case, we required a summary of reachable addresses on the stack. Abstract garbage collection combats merging in arguments by eliminating monotonicity for the abstract store; pushdown analysis eliminates the loss in return-flow precision by simulating the concrete call stack with a pushdown stack, thereby properly matching returns to call.

References

1. Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages*, pages 238–252. ACM Press, 1977.
2. Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In *POPL '79: Proceedings of the 6th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 269–282. ACM, 1979.
3. Saumya K. Debray and Todd A. Proebsting. Interprocedural control flow analysis of first-order programs with tail-call optimization. *ACM Trans. Program. Lang. Syst.*, 19(4):568–585, 1997.
4. Christopher Earl, Matthew Might, and David Van Horn. Pushdown control-flow analysis of higher-order programs. In *Proceedings of the 2009 Workshop on Scheme and Functional Programming*, Montreal, Quebec, Canada, August 2010.
5. Matthias Felleisen and Daniel P. Friedman. A calculus for assignments in higher-order languages. In *POPL '87: Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 314+. ACM, 1987.
6. Jan Midtgaard and Thomas P. Jensen. Control-flow analysis of function calls and returns by abstract interpretation. In *ICFP '09: Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming*, pages 287–298. ACM, 2009.
7. Matthew Might. *Environment Analysis of Higher-Order Languages*. PhD thesis, Georgia Institute of Technology, June 2007.
8. Matthew Might and Olin Shivers. Improving flow analyses via ΓCFA: Abstract garbage collection and counting. In *ICFP '06: Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming*, pages 13–25. ACM, 2006.
9. Thomas Reps, Susan Horwitz, and Mooly Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL '95: Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 49–61. ACM, 1995.
10. Olin G. Shivers. *Control-Flow Analysis of Higher-Order Languages*. PhD thesis, Carnegie Mellon University, 1991.
11. Michael Sipser. *Introduction to the Theory of Computation*. Course Technology, 2 edition, February 2005.
12. David Van Horn and Matthew Might. Abstracting abstract machines. In *ICFP '10: Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming*, September 2010.
13. Dimitrios Vardoulakis and Olin Shivers. CFA2: A context-free approach to control-flow analysis. In Andrew D. Gordon, editor, *Programming Languages and Systems*, volume 6012, chapter 30, pages 570–589. Springer Berlin Heidelberg, 2010.

A Classical control-flow analysis

This section presents traditional control-flow analysis for reference and comparison with pushdown and stack-summarizing control flow analysis. Classical control-flow analysis for ANF operates over the abstract state-space in Figure 2. Our classical formulation follows Van Horn and Might's technique of allocating abstract continuations in the store, as opposed to stacking them [12].

$$
\begin{array}{rll}
\hat{c} \in \widehat{Conf} = \widehat{State} \times \widehat{Addr} & & \text{[configurations]} \\
\hat{\varsigma} \in \widehat{State} = \mathsf{Exp} \times \widehat{Env} \times \widehat{Store} & & \text{[states]} \\
\hat{\rho} \in \widehat{Env} = \mathsf{Var} \rightharpoonup \widehat{Addr} & & \text{[environments]} \\
\hat{\sigma} \in \widehat{Store} = \widehat{Addr} \to \mathcal{P}\left(\widehat{Clo} \cup \widehat{Frame}\right) & & \text{[stores]} \\
\widehat{clo} \in \widehat{Clo} = \mathsf{Lam} \times \widehat{Env} & & \text{[closures]} \\
\hat{\phi} \in \widehat{Frame} = \mathsf{Var} \times \mathsf{Exp} \times \widehat{Env} \times \widehat{Addr} & & \text{[stack frames]} \\
\widehat{rp}, \hat{a} \in \widehat{Addr} \text{ is a } \textit{finite} \text{ set of addresses} & & \text{[addresses]}
\end{array}
$$

Fig. 2. Abstract configuration-space for classical control-flow analysis.

To complete the abstract semantics we need to define program injection, atomic expression evaluation, reachable configurations, transition relation, address allocation, and abstraction function:

Program injection The abstract injection function $\hat{\mathcal{I}} : \mathsf{Exp} \to \widehat{Conf}$ pairs an expression with an empty environment, an empty store and an empty stack to create the initial abstract configuration:

$$\hat{c}_0 = \hat{\mathcal{I}}(e) = (e, [], [], \widehat{null}),$$

where $\widehat{null}$ is an address bound to nothing in the store, thus representing an empty stack.

Atomic expression evaluation The abstract atomic expression evaluator, $\hat{\mathcal{A}} : \mathsf{Atom} \times \widehat{Env} \times \widehat{Store} \to \mathcal{P}(\widehat{Clo} \cup \widehat{Frame})$, returns the value of an atomic expression or a stack frame in the context of an environment and a store; note how it returns a set:

$$
\begin{array}{rll}
\hat{\mathcal{A}}(lam, \hat{\rho}, \hat{\sigma}) = \{(lam, \rho)\} & & \text{[closure creation]} \\
\hat{\mathcal{A}}(v, \hat{\rho}, \hat{\sigma}) = \hat{\sigma}(\hat{\rho}(v)) & & \text{[variable look-up].}
\end{array}
$$

Reachable configurations The abstract program evaluator $\hat{\mathcal{E}} : \mathsf{Exp} \to \mathcal{P}(\widehat{Conf})$ returns all of the configurations reachable from the initial configuration:

$$\hat{\mathcal{E}}(e) = \left\{ \hat{c} : \hat{\mathcal{I}}(e) \rightsquigarrow^* \hat{c} \right\}.$$

Transition relation The abstract transition relation $(\rightsquigarrow) \subseteq \widehat{Conf} \times \widehat{Conf}$ has three rules, two of which have become nondeterministic. A tail call may fork because there could be multiple abstract closures that it is invoking:

$$\begin{aligned}
(\overbrace{([\![\texttt{(}f\ \textit{æ}\texttt{)}]\!], \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, \widehat{rp}) \rightsquigarrow ((e, \hat{\rho}'', \hat{\sigma}'), \widehat{rp}), \text{ where} \\
([\![\texttt{(}\lambda\ \texttt{(}v\texttt{)}\ e\texttt{)}]\!], \hat{\rho}') \in \hat{\mathcal{A}}(f, \hat{\rho}, \hat{\sigma}) \\
\hat{a} = \widehat{alloc}(v, \hat{\varsigma}) \\
\hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}] \\
\hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathcal{A}}(æ, \hat{\rho}, \hat{\sigma})].
\end{aligned}$$

The partial order for stores is:

$$(\hat{\sigma} \sqcup \hat{\sigma}')(\hat{a}) = \hat{\sigma}(\hat{a}) \cup \hat{\sigma}'(\hat{a}).$$

A non-tail call builds a frame, adds it to the store, and evaluates the call:

$$\begin{aligned}
(\overbrace{([\![\texttt{(let ((}v\ \textit{call}\texttt{)) }e\texttt{)}]\!], \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, \widehat{rp}) \rightsquigarrow ((\textit{call}, \hat{\rho}, \hat{\sigma}), \widehat{rp}'), \text{ where} \\
\widehat{rp}' = \widehat{alloc}(v, \hat{\varsigma}) \\
\hat{\sigma}' = \hat{\sigma} \sqcup [\widehat{rp}' \mapsto (v, e, \hat{\rho}, \widehat{rp})].
\end{aligned}$$

A function return may fork because there could be multiple frames bound to the current return pointer:

$$\begin{aligned}
(\overbrace{(æ, \hat{\rho}, \hat{\sigma})}^{\hat{\varsigma}}, \widehat{rp}) \rightsquigarrow ((e, \hat{\rho}'', \hat{\sigma}'), \widehat{rp}'), \text{ where} \\
(v, e, \hat{\rho}', \widehat{rp}') \in \hat{\sigma}(\widehat{rp}) \\
\hat{a} = \widehat{alloc}(v, \hat{\varsigma}) \\
\hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}] \\
\hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathcal{A}}(æ, \hat{\rho}, \hat{\sigma})].
\end{aligned}$$

Allocation, polyvariance and context-sensitivity In the abstract semantics, the abstract allocation function $\widehat{alloc} : \mathsf{Var} \times \widehat{State} \to \widehat{Addr}$ determines the polyvariance of the analysis (and, by extension, its context-sensitivity). The abstract allocation function is overloaded to assign return pointers (addresses) to abstract stack frames: $\widehat{alloc} : \widehat{Frame} \times \widehat{State} \to \widehat{Addr}$. In a control-flow analysis, *polyvariance* literally refers to the number of abstract addresses (variants) there are for each variable.

Abstraction function The abstraction function (α) converts any structure from the concrete semantics (Figure 1) into an abstract form of the same structure (Figure 2). (The abstraction function is defined in Appendix B. While not specifically defined for these semantics, the abstraction function there can easily be modified to work with return pointers.)

Comparison to pushdown control-flow analysis The abstract semantics of pushdown control-flow analysis are similar to those of the abstract semantics for classical control-flow analysis (Figure 2). However, the few key differences are worth noting.

Foremost, we are not working with the configuration-space directly; rather we deal with the control-state-space. Hence, a configuration is now defined as a state and a stack paired together. The results of pushdown control-flow analysis are rooted pushdown systems rather than nondeterminisitic finite automata. A rooted pushdown system handles the stack and configurations implicitly, so we use the control-state-space instead of the configuration-space.

The next change is that the store only contains bindings from addresses to closures, instead of from addresses to closures and frames. The abstraction of the store creates imprecision. By keeping the frames out of the store (and in a precise structure) we avoid this imprecision for continuations.

The final difference is that frames no longer contain return pointers. Again, this is because the enriched abstract transition system encapsulates that information precisely.

B Abstraction Functions

This section defines the abstraction functions used throughout this paper. In particular, the following abstraction functions links the concrete and the abstract configuration-spaces in Section 3. The abstraction function recurs structurally:

$$\begin{aligned}
\alpha(\varsigma, \kappa) &= (\alpha_{State}(\varsigma), \alpha_{Kont}(\kappa)) && \text{[configuration abstraction]}\\
\alpha_{State}(e, \rho, \sigma) &= (e, \alpha_{Env}(\rho), \alpha_{Store}(\sigma)) && \text{[state abstraction]}\\
\alpha_{Env}(\rho)(v) &= \alpha_{Addr}(\rho(v)) && \text{[environment abstraction]}\\
\alpha_{Store}(\sigma)(\hat{a}) &= \bigsqcup_{\alpha_{Addr}(a)=\hat{a}} \alpha_{Clo}(\sigma(a)) && \text{[store abstraction]}\\
\alpha_{Clo}(lam, \rho) &= \{(lam, \alpha_{Env}(\rho))\} && \text{[closure abstraction]}\\
\alpha_{Kont}(\langle\phi_1, \ldots, \phi_n\rangle) &= \langle\alpha_{Frame}(\phi_1), \ldots, \alpha_{Frame}(\phi_n)\rangle && \text{[stack abstraction]}\\
\alpha_{Frame}(v, e, \rho) &= (v, e, \alpha_{Env}(\rho)) && \text{[frame abstraction].}
\end{aligned}$$

Just as address-allocation is a parameter, the address abstraction function, $\alpha_{Addr} : Addr \to \widehat{Addr}$, is a parameter for the abstract semantics.

For Sections 5, 6, and 7, the abstraction function for configurations is:

$$\alpha(\varsigma, \kappa) = (\alpha_{State}(\varsigma), \alpha_S(\kappa)) \qquad \text{[configuration abstraction],}$$

where the stack summarization function, α_S, is a parameter as described in Section 5.

C Building a Dyck configuration graph

C.1 Building a Dyck state graph for PDCFA

A fixed-point approach to building Dyck state graphs for PDCFA is best presented in [4]. The algorithm in Figure 3 is a similar iterative algorithm, but it is formulated for stack-summarizing control-flow analysis (Section 6). The underlying approaches are similar. In fact, for the algorithm in Figure 3, replacing configurations with states and switching the transition relation used throughout to the transition relation of Section 3 ($\curvearrowright$) is enough to convert the algorithm to build standard Dyck state graphs. The main difference is that the algorithm presented here examines a single state, transition, or shortcut edge each iteration, whereas the fixed-point algorithm examines the entire frontier each iteration.

C.2 Building a Dyck state graph for SSCFA

The algorithm in Figure 3 builds the Dyck configuration graph and the ϵ-closure graph for a given program e. It uses the worklists, ΔS, ΔE, and ΔH, to maintain the frontier of unexplored configurations, transitions, and shortcut edges respectively. The while loop runs until all the worklists are empty, which is exactly when everything reachable has been explored. Each iteration of the while loop, explores one previously unexplored shortcut edge, transition, or configuration.

Each new configuration, transition, and shortcut edge can imply other configurations, transitions, and shortcut edges globally. Thus, after each configuration, transition, or shortcut edge is explored, the implied configurations, transitions, and shortcut edges are added to the worklists.

The procedure *addShort* finds all the configurations, transitions, and shortcut edges implied by the given shortcut edge. The only new transitions implied by a shortcut edge are pop transitions that are enabled by a new push transition. The only new configurations are those in the newly implied pop transitions.

The procedure *addEdge* finds all the configurations, transitions, and shortcut edges implied by the given transition. There are three types of transitions: First, there are no-op (ϵ) transitions, which immediately become shortcut edges, and so are added to the ϵ-closure graph and are expanded. Next, there are push transitions, which imply new pop transitions (and configurations from these pop transitions) as well as new shortcut edges (from pre-existing pop transitions). Finally, there are pop transitions, which imply only new new shortcut edges from pre-existing push transitions.

The last procedure *Explore* finds all the configurations, transitions, and shortcut edges implied by the given configuration. A configuration cannot imply any shortcut edges directly. However, a configuration can imply new no-op, push, or pop transitions as well as new configurations from these transitions.

procedure *BuildDyck*(e)
 $\hat{c}_0 \leftarrow \hat{\mathcal{I}}_{\hat{c}}(e)$; $G \leftarrow (\emptyset, \Gamma, \emptyset, \hat{c}_0)$; $G_\epsilon \leftarrow (\emptyset, \emptyset)$; $\Delta S \leftarrow \{\hat{c}_0\}$; $\Delta E \leftarrow \emptyset$; $\Delta H \leftarrow \emptyset$
 while ($\Delta S \neq \emptyset$ **or** $\Delta E \neq \emptyset$ **or** $\Delta H \neq \emptyset$)
 if ($\Delta H \neq \emptyset$), **let** $(\hat{c}, \hat{c}') \in \Delta H$
 $(\Delta S', \Delta E', \Delta H') \leftarrow addShort(G, G_\epsilon)(\hat{c} \rightarrowtail \hat{c}')$
 $(S, H) \leftarrow G_\epsilon$
 $G_\epsilon \leftarrow (S, H \cup \{(\hat{c}, \hat{c}')\})$
 $(\Delta S, \Delta E, \Delta H) \leftarrow (\Delta S \cup \Delta S', \Delta E \cup \Delta E', \Delta H \cup \Delta H' - \{(\hat{c}, \hat{c}')\})$
 else if ($\Delta E \neq \emptyset$), **let** $(\hat{c}, g, \hat{c}') \in \Delta E$
 $(\Delta S', \Delta E', \Delta H') \leftarrow addEdge(G, G_\epsilon)(\hat{c} \rightarrowtail^g \hat{c}')$
 $(S, \Gamma, E, \hat{c}_0) \leftarrow G$
 $G \leftarrow (S, \Gamma, E \cup \{(\hat{c}, g, \hat{c}')\}, \hat{c}_0)$
 $(\Delta S, \Delta E, \Delta H) \leftarrow (\Delta S \cup \Delta S', \Delta E \cup \Delta E' - \{(\hat{c}, g, \hat{c}')\}, \Delta H \cup \Delta H')$
 else if ($\Delta S \neq \emptyset$), **let** $\hat{c} \in \Delta S$
 $(\Delta S', \Delta E', \Delta H') \leftarrow Explore(G, G_\epsilon)(\hat{c})$
 $(S, \Gamma, E, \hat{c}_0) \leftarrow G$
 $(S, H) \leftarrow G_\epsilon$
 $(G, G_\epsilon) \leftarrow ((S \cup \{\hat{c}\}, \Gamma, E, \hat{c}_0), (S \cup \{\hat{c}\}, H))$
 $(\Delta S, \Delta E, \Delta H) \leftarrow (\Delta S \cup \Delta S' - \{\hat{c}\}, \Delta E \cup \Delta E', \Delta H \cup \Delta H')$
 return G, G_ϵ

procedure $addShort(G, G_\epsilon)(\hat{c}, \hat{c}')$
 $(S, \Gamma, E, \hat{c}_0) \leftarrow G$; $(S, H) \leftarrow G_\epsilon$
 $\Delta E \leftarrow \left\{(\hat{c}', \hat{\phi}_-, \hat{c}_2) : (\hat{c}_1, \hat{\phi}_+, \hat{c}) \in E \text{ and } \hat{c}' \twoheadrightarrow^{\hat{\phi}_-} \hat{c}_2\right\}$
 $\Delta S \leftarrow \left\{\hat{c}_2 : (\hat{c}_1, \hat{\phi}_-, \hat{c}_2) \in \Delta E\right\}$
 $\Delta H \leftarrow \{(\hat{c}_1, \hat{c}') : (\hat{c}_1, \hat{c}) \in H\}$
 $\cup \{(\hat{c}, \hat{c}_2) : (\hat{c}', \hat{c}_2) \in H\}$
 $\cup \{(\hat{c}_1, \hat{c}_2) : (\hat{c}_1, \hat{c}), (\hat{c}', \hat{c}_2) \in H\}$
 return $\Delta S - S,\ \Delta E - E,\ \Delta H - H$

procedure $addEdge(G, G_\epsilon)(\hat{c} \rightarrowtail^g \hat{c}')$
 $(S, \Gamma, E, \hat{c}_0) \leftarrow G$; $(S, H) \leftarrow G_\epsilon$
 if $(g = \epsilon)$ **return** $addShort(G, (S, H \cup \{(\hat{c}, \hat{c}')\}))(\hat{c}, \hat{c}')$
 else if $(g = \hat{\phi}_+)$
 $\Delta E \leftarrow \left\{\hat{c}_1 \rightarrowtail^{\hat{\phi}_-} \hat{c}_2 : \hat{c}' \rightarrowtail \hat{c}_1 \in H \text{ and } \hat{c}_1 \twoheadrightarrow^{\hat{\phi}_-} \hat{c}_2\right\}$
 $\Delta S \leftarrow \left\{\hat{c}_2 : \hat{c}_1 \rightarrowtail^{\hat{\phi}_-} \hat{c}_2 \in \Delta E\right\}$
 $\Delta H \leftarrow \left\{\hat{c} \rightarrowtail \hat{c}_2 : \hat{c}' \rightarrowtail \hat{c}_1 \in H \text{ and } \hat{c}_1 \rightarrowtail^{\hat{\phi}_-} \hat{c}_2 \in E\right\}$
 return $\Delta S - S,\ \Delta E - E,\ \Delta H - H$
 else if $(g = \hat{\phi}_-)$
 return $\emptyset,\ \emptyset,\ \left\{\hat{c}_1 \rightarrowtail \hat{c}' : \hat{c}_2 \rightarrowtail \hat{c} \in H \text{ and } \hat{c}_1 \rightarrowtail^{\hat{\phi}_+} \hat{c}_2 \in E\right\} - H$

procedure $Explore(G, G_\epsilon)(\hat{c})$
 $(S, \Gamma, E, \hat{c}_0) \leftarrow G$; $(S, H) \leftarrow G_\epsilon$
 return $\{\hat{c}' : \hat{c} \twoheadrightarrow^g \hat{c}'\} - S,\ \{(\hat{c}, g, \hat{c}') : \hat{c} \twoheadrightarrow^g \hat{c}'\} - E,\ \emptyset$

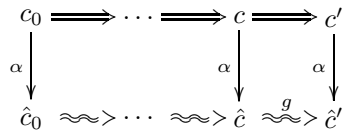
Fig. 3. Algorithm to build a Dyck configuration graph. Procedures *addShort*, *addEdge* and *Explore* determine what configurations, transitions, and shortcut edges are implied by a given shortcut edge, transition, or configuration, respectively.

C.3 Building a Dyck state graph for SS$\mathit{\Gamma}$CFA

Finally, we let the procedures of Figure 3 use this transition relation ($\hat{\approx}\!\!>_{AGC}$) and the reachable address push function $\mathbf{push}^{ra}$. Now the procedure BUILDDYCK of Figure 3 computes stack-summarizing control-flow analysis with abstract garbage collection soundly.

D Proofs

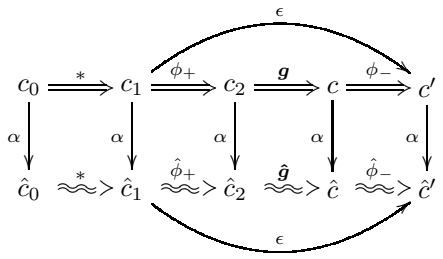
Proof (of Theorem 1). Without loss of generality, assume that the path from the initial configuration c_0 to c is length n, so the path to the new configuration c' is $n+1$ transitions from the initial configuration. The inductive hypothesis is that the theorem holds for all paths of length less than or equal to n. So far the scenario can be diagrammed as below:



We now have three cases depending on what g is:

- $\hat{c} \hat{\approx}\!\!>^{\epsilon} \hat{c}' = (\hat{\varsigma}'', \widehat{ss}')$
 No change is made to the stack in the concrete, thus the stacks are equal: $\kappa = \kappa'$. Likewise, no change is made in the abstract, thus: $\widehat{ss} = \widehat{ss}'$. Since the first stack is subsumed by the first stack summary, the second stack must be subsumed by the second stack summary: $\alpha_S(\kappa') \sqsubseteq_S \widehat{ss}'$.
- $\hat{c} \hat{\approx}\!\!>^{\hat{\phi}+} \hat{c}' = (\hat{\varsigma}'', \widehat{ss}')$
 A frame ϕ, such that $\alpha(\phi) \sqsubseteq \hat{\phi}$, must be pushed onto the stack in the concrete, so the stacks are related thusly: $\phi : \kappa = \kappa'$. Likewise, the stack summaries are so related: $\mathbf{push}(\hat{\phi}, \widehat{ss}) \sqsubseteq_S \widehat{ss}'$. By the constraint on all push operations: $\alpha_S(\phi : \kappa) \sqsubseteq_S \mathbf{push}(\hat{\phi}, \alpha_S(\kappa))$. We can make the following replacements: $\alpha_S(\kappa') \sqsubseteq_S \mathbf{push}(\hat{\phi}, \widehat{ss})$. By the definition of the transitive relation ($\hat{\approx}\!\!>$): $\widehat{ss}' = \mathbf{push}(\hat{\phi}, \widehat{ss})$. Finally we have: $\alpha_S(\kappa') \sqsubseteq_S \widehat{ss}'$.
- $\hat{c} \hat{\approx}\!\!>^{\hat{\phi}-} \hat{c}' = (\hat{\varsigma}'', \widehat{ss}')$
 A frame ϕ, such that $\alpha(\phi) \sqsubseteq \hat{\phi}$, must be popped from the stack in the concrete, so the stacks are related like this: $\kappa = \phi : \kappa'$. We know that configuration c is reachable from the initial configuration, which has an empty stack. Since the transition from the configuration pops off a frame ϕ, it does not currently have an empty stack. So there exists a path, $c_0 \Rightarrow^* c_1 \Rightarrow^{\phi+} c_2 \Rightarrow^{\boldsymbol{g}} c$, such that the net of the stack actions after the push, $\lfloor \boldsymbol{g} \rfloor$, is empty. Let $c_1 = (\varsigma_1, \kappa_1)$. Since the net of the stack actions is empty, the stack at the configuration before the last previously unmatched push, κ_1, is identical to the stack after the current pop, κ': $\kappa_1 = \kappa'$.

By the inductive hypothesis, there is a path through the Dyck configuration graph that parallels and mimics the path above. Thus, there is a configuration $\hat{c}_1 = (\hat{\varsigma}_1, \widehat{ss}_1)$ such that $\alpha(c_1) \sqsubseteq \hat{c}_1$. Also by this path, there is a sub-path from the configuration $\hat{c}_1$ to the new configuration $\hat{c}'$ that makes no changes to the stack. Therefore, there is a shortcut edge between these two configurations $\hat{c}_1$ and $\hat{c}'$. The proof of the first case for shortcut edges works for no-op transitions. Thus the stack summaries at these two configurations are the same: $\widehat{ss}_1 = \widehat{ss}'$.
The current situation is as follows:



Since the configuration before the last previously unmatched push, c_1 is subsumed by its equivalent in the Dyck configuration graph, $\hat{c}_1$, its stack is subsumed by the stack summary of configuration $\hat{c}_1$: $\alpha_S(\kappa_1) \sqsubseteq_S \widehat{ss}_1$. Since this stack and this stack summary are identical to the stack κ' and the stack summary $\widehat{ss}'$ respectively, we have: $\alpha_S(\kappa') \sqsubseteq_S \widehat{ss}'$. $\square$